

Two-dimensional traffic rules and the density classification problem

Nazim Fatès

Inria Nancy - Grand Est LORIA UMR 7503

`nazim.fates@loria.fr`

Irène Marcovici

Institut Élie Cartan de Lorraine, Université de Lorraine

`irene.marcovici@univ-lorraine.fr`

Siamak Taati

Mathematics Institute, Leiden University

`siamak.taati@gmail.com`

April 18, 2016

Abstract

The density classification problem is the computational problem of finding the majority in a given array of votes, in a distributed fashion. It is known that no cellular automaton rule with binary alphabet can solve the density classification problem. On the other hand, it was shown that a probabilistic mixture of the traffic rule and the majority rule solves the one-dimensional problem correctly with a probability arbitrarily close to one. We investigate the possibility of a similar approach in two dimensions. We show that in two dimensions, the particle spacing problem, which is solved in one dimension by the traffic rule, has no cellular automaton solution. However, we propose exact and randomized solutions via interacting particle systems. We assess the performance of our models using numeric simulations.

1 Introduction

Let us imagine a medium composed of a great number of cells arranged regularly on a grid. Each cell is linked with its immediate neighbours and the only thing it can do is to change its own state according to the state of its neighbours. Can we compute with such a medium? And what happens if the updates occur at random times? And what if the cells are subject to noise?

In order to study this robustness mechanisms on a mathematical basis, we will here focus on two simple computational problems. The first problem is the *density classification problem*, which is the problem of finding the majority state in a distributed fashion. In the original setting, the computational protocol is required to be local and parallel, and use no extra memory other than the evolving configuration itself. The protocol is also required to be scalable, which means it must perform the task on configurations of arbitrary size. In other words, we look for a cellular automaton rule that performs the task. In this paper, we also consider variants of this problem in which the process is allowed to be asynchronous, non-deterministic or random.

This problem has attracted a considerable amount of attention these last years. It is trivial in most settings but it is not easy to solve in the case of cellular automata. The difficulty comes from the necessity to reach a consensus on the state of the cells: the system should converge to a situation with all 1s or all 0s, depending on whether the initial state contains more 0s or more 1s, respectively.

Inspired by the work of Gács, Kurdiumov and Levin, in 1988, Packard formulated this problem as a challenge to study genetic algorithms [10]. This triggered a wide competition to find rules with an increasing quality of classification. In 1995, Land and Belew proved that no perfect solution exists for one-dimensional deterministic systems [8]. Recently, this fact was re-demonstrated with a simpler argument and the proof was extended to probabilistic rules and to any dimension [2]. It was even shown that for any candidate solution there are configurations with a density close to 0 and 1 that are misclassified [7].

Since then, different variants of the problem have been proposed and it has been shown by various authors that relaxing one of the conditions of the problem is often sufficient to find perfect solutions [3]. In particular, Fuk s proposed to combine two rules sequentially to obtain a perfect solution, see Ref. [6] and references therein. Probabilistic cellular automata could provide another interesting framework: it was discovered that although no perfect rule exists, it is possible to find a family of one-dimensional nearest-neighbour rules for which the probability of making an error of classification can be made as small as wanted [5]. The perfect solution can thus be approximated – but not reached ! – at the cost of an increase in the average time to reach a consensus.

The construction proposed for building this family of rules consists of mixing stochastically two well-known rules: the traffic rule, which introduces space between particles, and the majority rule, which has a “homogenising” effect. In this text, we ask whether there also exist a “close-to-perfect” solution for two-dimensional cellular automata. At first sight, one does not see why the problem should be significantly different for two-dimensional systems. However, there is no such thing as a “traffic” rule in two dimensions (2D). If we decompose a 2D grid in layers and apply a classical traffic rule on each layer, then different consensus might be attained and there is no obvious means on how one can obtain the “right” global consensus from a collection of local consensus.

We call the problem that is solved by the traffic rule in one dimension the *particle spacing problem*. We tackle this problem in two dimensions. We then




Figure 1: Transformation of a random initial condition. Under some conditions (see Sec. 4 p. 12), the system will most probably converge to $1^{\mathcal{L}}$, which is a correct classification since the initial density is greater than $1/2$.

(partially) solve the density classification problem by combining our particle spacing model with a local majority rule.

The outline of the article is as follows. After presenting the basic definitions and properties of our models in Sec. 2, we show the advantage of using interacting particle systems to tackle the problem (Sec. 3). We then present a concrete solution and analyse its behaviour with numerical simulations in Sec. 4.

2 Basics

2.1 Setting

In dimension $d \geq 1$, we set the cellular space to be a grid with periodic boundary conditions, defined by $\mathcal{L} = (\mathbb{Z}/n_1\mathbb{Z}) \times \cdots \times (\mathbb{Z}/n_d\mathbb{Z})$, for some $n_1, \dots, n_d \geq 1$. The number of cells of \mathcal{L} is $N_{\mathcal{L}} = n_1 \cdots n_d$. We say that the grid is *even-sized* if n_1, \dots, n_d are all even.

Each cell of this space can hold a binary state, so that the set of states is denoted by $Q = \{0, 1\}$.

The set of configurations is denoted by $\mathcal{E} = \{0, 1\}^{\mathcal{L}}$.

For a configuration $x \in \mathcal{E}$, and a state $q \in Q$, we define the density of state q by: $d_q(x) = \frac{1}{N_{\mathcal{L}}} \text{Card}\{i \in \mathcal{L}; x_i = q\}$.

For a given configuration $x \in \mathcal{E}$, we say that cell i is *isolated* if none of its adjacent cells is in state x_i .

For $q \in Q$, we say that a configuration $x \in \mathcal{E}$ is a *q-archipelago* if all the cells in state q are isolated, i.e., if x does not contain two adjacent cells in state q . We denote by \mathcal{A}_q the set of q -archipelagos. In particular, if $x \in \mathcal{A}_q$, then $d_q(x) \leq 1/2$.

We also introduce $\mathcal{A} = \mathcal{A}_0 \cup \mathcal{A}_1$, the set of all archipelagos.

2.2 Presentation of the problem

Recall that in this paper, we study two computational problems. The *density classification task* is the task of transforming a given configuration $x \in Q^{\mathcal{L}}$ into one of the two uniform configurations $1^{\mathcal{L}}$ or $0^{\mathcal{L}}$ depending on which of 1 or 0 has strict majority in x (see Fig. 1). More specifically, given an input $x \in Q^{\mathcal{L}}$,




Figure 2: Transformation of a random initial configuration into an archipelago.

a computational process performing the density classification task must return $1^{\mathcal{L}}$ if $d_1(x) > 1/2$ and $0^{\mathcal{L}}$ if $d_0(x) > 1/2$. (Generally, the case $d_1(x) = d_0(x)$ is avoided.)

Our approach to solve the density classification problem is via another problem which we call the particle spacing problem. The *particle spacing problem* is the computational problem of rearranging the “particles” (say, symbols 1) on a configuration $x \in Q^{\mathcal{L}}$ so as to obtain an archipelago configuration (see Fig. 2). Again we require the computational process to be local and scalable, but we also require it to be *conservative*: at every step of the process, the number of particles (symbols 1) must be preserved.

There are two possible variants for the latter problem. In the *strict* spacing problem, we require that the sets \mathcal{A}_1 and \mathcal{A}_0 are *absorbing*, in the sense that as soon as the process enters \mathcal{A}_q (for either $q = 1$ or $q = 0$), it cannot leave it. In the *loose* variant of the problem, we require the process to eventually remain in \mathcal{A}_q . Note that the latter is equivalent to the condition that the computation reaches an absorbing subset of \mathcal{A}_1 or \mathcal{A}_0 .

For both problems, our purpose is to build a solution with a cellular system; this means that we have a set of interacting components, which can have a deterministic or stochastic behaviour, and interact only locally. We will here consider cellular automata (deterministic or stochastic) and interacting particle systems.

2.3 Known results with cellular automata

We now present the principal known results concerning solutions of the density classification problem and of the particle spacing problem, using cellular automata.

A cellular automaton $F : \mathcal{E} \rightarrow \mathcal{E}$ is defined by a neighbourhood $\mathcal{N} = (v_1, \dots, v_k) \in \mathcal{L}^k$, and by a local rule $f : Q^k \rightarrow Q$, which defines the global rule F , mapping a configuration x to the configuration $F(x)$ defined by:

$$\forall c \in \mathcal{L}, F(x)_c = f(x_{c+v_1}, \dots, x_{c+v_k}).$$

Proposition 1. *For any $d \geq 1$, there is no deterministic cellular automaton solving the density classification problem. For $d = 1$, this means that there is no local rule f such that for any $n \geq 1$, and any $x \in \{0, 1\}^{\mathbb{Z}/n\mathbb{Z}}$,*

$$d_q(x) > 1/2 \implies \exists T \geq 0, \forall t \geq T, F^t(x) = q^{\mathbb{Z}/n\mathbb{Z}}.$$

The first proof was given by Land and Belew in 1995 for dimension $d = 1$ [8]. A simplified proof was proposed in 2013 for any dimension $d \geq 1$ [2]. These results apply to deterministic cellular automata. One can ask whether stochastic transition rules could help to solve the problem. Using the same argument as for deterministic cellular automata, one can prove that there are no probabilistic cellular automata solving perfectly the density classification problem [2].

However, for $d = 1$, Fatès has provided a family of probabilistic cellular automata solving the density classification problem with an arbitrary precision [5]. This means that the probability of making a bad classification can be reduced to as low as necessary, at the cost of an increase of the average time of convergence to the uniform configuration.

The family of rules is defined with a real parameter $\epsilon > 0$. The local rule consists at each time step, for each cell independently, in applying the traffic rule with probability $1 - \epsilon$ and the majority rule with probability ϵ . The traffic rule (rule 184 with Wolfram's notations) is a conservative rule, which moves the 1s to the right whenever possible. It has a spacing effect. The majority rule allows the convergence to the uniform fix point, once particles have been spaced.

In order to extend this result to higher dimensions, one would like to design a rule having the same behaviour as the traffic rule, that is, to be able to compose a rule that solves the spacing problem with a majority rule. Unfortunately, this is not possible in the classical framework of cellular automata.

Proposition 2. *1. In dimension 1, the traffic cellular automaton F_{184} solves the spacing problem. Precisely, it satisfies: for all $n \geq 1$ and all $x \in \{0, 1\}^{\mathbb{Z}/n\mathbb{Z}}$,*

$$d_q(x) \leq 1/2 \implies \forall t \geq n/2, F_{184}^t(x) \in \mathcal{A}_q.$$

Furthermore, $\forall q \in \{0, 1\}, F_{184}(\mathcal{A}_q) \subset \mathcal{A}_q$, so F_{184} is a solution to the strict spacing problem.

2. In dimension $d \geq 2$, there are no cellular automata that solve the spacing problem.

Proof. The fact that the traffic cellular automaton spaces configurations is a “folk” result.

Let now F be a d -dimensional deterministic cellular automaton. If $x \in \mathcal{E}$ is a configuration with a symmetry of translation, then the symmetry is conserved by the evolution of the automaton. Formally, if there exists $\delta \in \mathbb{Z}^d$ such that $\forall c \in \mathcal{L}, x_c = x_{c+\delta}$, then $\forall t \in \mathbb{N}, \forall c \in \mathcal{L}, F^t(x)_c = F^t(x)_{c+\delta}$. As a consequence, deterministic cellular automata can not solve the particle spacing problem in dimension $d \geq 2$. To see why, simply consider a configuration with all 0s, except one line which is made of cells with all 1s: if the rule is conservative, this line can not disappear. \square

By its very nature, a “truly” probabilistic rule can not solve the particle spacing problem. Indeed, as soon as there exists a configuration for which one cell has a non-deterministic outcome, we cannot ensure that the number of particles will be preserved. We leave open the question as to whether there

exists a probabilistic rule which would solve the density classification problem with an arbitrary precision, in dimension 2 or more.

3 Particle systems solutions to the spacing problem

We have seen that deterministic cellular automata are in some sense too rigid to allow us to solve the spacing problem on grids, because they do not allow to break translation symmetries. On the other hand, probabilistic cellular automata can break these symmetries, but they do not allow an exact conservation of the number of particles.

We now propose to combine the strength of both models with interacting particle systems: we update cells by *pairs*, which allows conservation of particles, and the pairs are chosen *randomly*, which allows us to break symmetries. The effect of the local rule is to exchange the cell's states or to leave them unchanged, depending on the states of the neighbouring cells of the pair.

Let us formalize the definition of the *interacting particle systems* (IPS) we consider. From now on, we will consider two-dimensional grids. Note that most results can be adapted to higher-dimensional lattices.

3.1 Our model of IPS

Let \mathcal{N}_i and \mathcal{N}_p be two finite tuples of \mathbb{Z}^2 , corresponding to the *interaction neighbourhood* and the *perception neighbourhood*.

We define the set of interacting pairs by

$$\mathcal{I} = \{(c, c + \delta) ; c \in \mathcal{L}, \delta \in \mathcal{N}_i\}.$$

The global rule is a function $\Phi : \mathcal{E} \times \mathcal{I} \rightarrow \mathcal{E}$. It takes in argument a configuration and a pair of cells to update, and maps it to the configuration that represents the next state of the system. The image $y = \Phi(x, (c, c'))$ is defined by:

$$(y_c, y_{c'}) = \phi((x_{c+k}, k \in \mathcal{N}_p), (x_{c'+k}, k \in \mathcal{N}_p)) \text{ and for } d \notin \{c, c'\}, y_d = x_d,$$

where $\phi : \{0, 1\}^{\mathcal{N}_p} \times \{0, 1\}^{\mathcal{N}_p} \rightarrow \{0, 1\}^2$ is the local rule that gives the new states of the pair of cells as a function of the states of their perception neighbourhood.

This rule is conservative if the image y always satisfies $(y_c, y_{c'}) = (x_c, x_{c'})$ or $(y_c, y_{c'}) = (x_{c'}, x_c)$.

Let $(u_t)_{t \in \mathbb{N}} \in \mathcal{I}^{\mathbb{N}}$ be a sequence of interacting pairs (in the following, the u_t are chosen uniformly at random independently in \mathcal{I}). Starting from an initial condition $x \in \mathcal{E}$, the system will evolve according to the sequence of states (or orbit) $(x^t)_{t \geq 0}$ defined by $x^0 = x$ and $x^{t+1} = \Phi(x^t, u_t)$ for any $t \geq 0$.

Given an IPS rule Φ , we say that a set $A \subset \mathcal{E}$ is an *absorbing set* if: $\forall x \in A, \forall u \in \mathcal{I}, \Phi(x, u) \in A$. We say that A is *reachable from any configuration* if:

$$\forall x^0 \in A, \exists T \geq 0, \exists (u_t)_{1 \leq t \leq T} \in \mathcal{I}^T, x^T \in A.$$

We say that A is a *sink* if it is an absorbing set, which is reachable from any configuration.

In terms of IPS, we say that Φ is a solution to the strict spacing problem if it is a conservative IPS such that the set \mathcal{A} of archipelagos is a sink.

3.2 No solution to the strict spacing problem

In order to solve the spacing problem, an idea is to design a rule, such that its evolution would result in decreasing the *energy* of the configuration, that is, the number of adjacent cells in same state. This idea will be used to propose an approximate solution in Section 3.4. However, the next proposition proves that this idea does not allow us to solve the strict spacing problem. This is due to the existence of configurations that are not archipelagos but for which each cell can “believe” that it is part of an archipelago (by looking at the cells located within a finite range), see Fig. 3.

Proposition 3. *There is no IPS solution to the strict spacing problem.*

Proof. By contradiction, assume that we have an IPS rule that is a solution to the strict spacing problem. For $c \in \mathbb{Z}^2$, and $r \geq 1$, let us introduce the notation $B(c, r) = \{c + \delta; \|\delta\|_1 \leq r\}$.

The interaction and the perception neighbourhoods being bounded, there exist $r \geq 1$ such that $\mathcal{N}_p \subset B(0, r)$, and $\Delta \geq 1$ such that for any $(c, c') \in \mathcal{I}$, the set $B(c, r) \cup B(c', r)$ has a diameter smaller than Δ .

Let us consider the lattice $\mathcal{L} = (\mathbb{Z}/n\mathbb{Z})^2$, with $n = 4k$, for some $k > \Delta$. We define the configuration $x \in \mathcal{E}$, by: $x_{i,j} = 0$ if and only if $(j-i) \in \{0\} \cup \{2i+1; 0 \leq i \leq k-1\} \cup \{2k+2i; 1 \leq i \leq k-1\}$. This configuration is represented on Figure 3 for $k = 4$. It is not an archipelago. Moreover, observe that every 1 has at least two adjacent cells in state 0, and the other way round. Note that this configuration has the same density of 0s and 1s, but it is also possible to build other counter-examples with various densities.

Let $(c, c') \in \mathcal{I}^2$ be such that $x_c \neq x_{c'}$. By construction of x , since $k > \Delta$, one can check that on the set of cells $B(c, r) \cup B(c', r)$, either there are no pairs of adjacent 1s or there are no pairs of adjacent 0s. Let us for example assume that there are no pairs of adjacent 1s.

We now consider the configuration y that coincides with x on the set $B(c, r) \cup B(c', r)$ and for which all the cells outside this set are in state 0. The configuration y is a 1-archipelago. For configuration y , the exchange between the states of cells c and c' is forbidden, since it would create a pair of adjacent 1s (because among c and c' , the cell in state 1 has at least two adjacent cells in state 0).

Let us go back to x . Since x and y are locally the same, the exchange of c and c' is forbidden. This implies that configuration x is a fixed point of the dynamics, which is a contradiction, because it is not an archipelago.

□

3.3 An IPS that synchronises checkerboards

We have seen that there is no IPS solution ensuring that once we have reached any archipelago configuration, we will remain in the set of archipelago configurations. At first sight, this could seem that there exist no solution to the spacing problem at all. However, to our surprise, we could notice that the loose problem is solvable. In fact, the *loose* problem is more demanding on the set of configurations that we do not leave. In other words, there might exist a subset \mathcal{A}' of archipelagos such that once the configuration reaches \mathcal{A}' , it remains in it.

It can be observed that in Figure 3, the problem comes from the fact that there are two checkerboards of different phases. If we were able to synchronize these two checkerboards, we would reach one of the two perfect checkerboards (since both 0 and 1 have density 1/2).

For $q \in \{0, 1\}$, we denote:

$$\begin{aligned} C_e^q &= \{x \in \mathcal{E} ; x_{i,j} = q \implies i + j \text{ is even}\}, \\ C_o^q &= \{x \in \mathcal{E} ; x_{i,j} = q \implies i + j \text{ is odd}\}. \end{aligned}$$

These sets correspond to “sub-checkerboards” in the sense that in C_e^q (resp. C_o^q), state q is the minority state and only appears on even (resp. odd) cells. We also introduce $C = C_e^0 \cup C_e^1 \cup C_o^0 \cup C_o^1$. It is a subset of \mathcal{A} .

Proposition 4. *There is an IPS solution to the loose spacing problem for an even-sized grid. Indeed, the rule Φ_C defined below is a conservative IPS having the property that the set C is a sink.*

To define Φ_C with a local description, we introduce its interaction neighbourhood $\mathcal{N}_i = \{-1, 0, 1\}^2 \setminus \{(0, 0)\}$, that is, we allow interactions between a cell and its eight nearest-neighbours. The perception neighbourhood corresponds to von Neumann neighbourhood, that is, $\mathcal{N}_p = \{(0, 0), (0, 1), (-1, 0), (0, -1), (1, 0)\}$.

Let \mathcal{I}_4 be the set of pairs of adjacent cells and \mathcal{D}_4 be the set of diagonal pairs. The interaction set is $\mathcal{I} = \mathcal{I}_4 \cup \mathcal{D}_4$. The rule will act differently on diagonal pairs on the one hand, and horizontal and vertical pairs on the other hand.

For $(i, j) \in \mathcal{I}$, let $\tau_{(i,j)} : \mathcal{E} \rightarrow \mathcal{E}$ be the function that exchanges the states x_i and x_j in configuration x . Precisely, $\tau(x, (i, j))$ is defined with:

$$\tau(x, (i, j))_k = \begin{cases} x_j & \text{if } k = i \\ x_i & \text{if } k = j \\ x_k & \text{if } k \notin \{i, j\}. \end{cases}$$

For a pair $(i, j) \in \mathcal{I}_4$, we define

$$\Phi_C(x, (i, j)) = \begin{cases} \tau(x, (i, j)) & \text{if both cells } i \text{ and } j \text{ are not isolated,} \\ x & \text{otherwise.} \end{cases}$$

For a pair $(i, j) \in \mathcal{D}_4$, we define $\Phi_C(x, (i, j)) = \tau(x, (i, j))$.

Note that if $x_i = x_j$, both cases above result in leaving x unchanged.

In practice, numerical simulations show that in order to improve the speed of convergence to the archipelago, it is more appropriate to do the exchanges with different probability rates, depending on the state of the neighbourhood of the pair. However, these parameter do not affect the reachability properties, this is why we will here work with the simplest version of the model.




Figure 3: Example of a configuration that locally looks like an archipelago.

Proposition 4. It is clear that the IPS Φ_C constructed above is conservative. We show that the set C of sub-checkerboards is a sink for Φ_C .

First, note that the set C is absorbing. Indeed, in every element of $C_e^q \cup C_o^q$, every q is isolated, and as a result $\Phi_C(x, (k, k')) = x$ for each $x \in C$ and $(k, k') \in \mathcal{I}_4$ (i.e., the horizontal/vertical rules keep x unchanged). On the other hand, for each $x \in C$ and $(k, k') \in \mathcal{D}_4$, clearly we have $\Phi_C(x, (k, k')) \in C$ (i.e., C is invariant under the diagonal rule). It remains to verify that C is reachable from any configuration.

Let $x \in \mathcal{E}$ be a configuration, and without loss of generality assume that $d_1(x) \leq 1/2$. Let us denote the number of even cells (i, j) with $x_{i,j} = 1$ by m_e and the number of odd cells (i, j) with $x_{i,j} = 1$ by m_o . If $m_e = 0$ or $m_o = 0$, then x is already in C , and there is nothing to prove. So, assume m_e and m_o are both non-zero. We describe a path from x to a configuration y that has $m_e + 1$ even cells in state 1. The claim would then follow by induction.

Let $M = \{0, 1\}^2 \subseteq \mathcal{L}$ be a window consisting of a 2×2 square. We first show that there is a position c such that the 2×2 pattern $(x_{c+k} : k \in M)$ contains either a single 1 or two adjacent 1s and two adjacent 0s (see Fig. ??). Indeed, there are three possibilities:

1. Every 2×2 window has at least 2 occurrences of 1. In this case, every 2×2 window must have precisely two occurrences of 1, for otherwise $d_1(x) > 1/2$. If every 2×2 window is "checkered", then x must be in

$C_e^1 \cup C_o^1$, which contradicts the assumption. Otherwise, x has a 2×2 pattern with two adjacent 1s and two adjacent 0s.

2. There is a 2×2 window with a single 1.
3. There is a 2×2 window with no 1s. Let us call such a window *empty*. Note that every window that either horizontally or vertically overlaps an empty window is either itself empty or has one of the two desired forms. Since we have assumed that x has at least one 1 ($m_e, m_o > 0$), not every 2×2 window can be empty. Consider a path from an empty window to a non-empty window, obtained by vertical and/or horizontal moves. The first time such a path reaches a non-empty window, we see a pattern with one of the two desired forms.

Let us consider a 2×2 window $c + M$ that has a single 1. Without loss of generality, we can assume that this single 1 is on an even cell. Since $m_o > 0$, there is at least one 1 somewhere outside $c + M$ on an odd cell. By applying the diagonal rule, we can move this single 1 to $c + M$, without changing any other cell, hence obtaining a new configuration that differs from x in exactly two cells. In this new configuration, the window $c + M$ has two adjacent 1s and two adjacent 0s.

Consider finally a 2×2 window $c + M$ with two adjacent 1s and two adjacent 0s. Let k be the odd cell in $c + M$ that has a 1 and k' be its adjacent cell in $c + M$ having state 0. Applying the interaction rule on pair (k, k') we end up in a configuration y with one more even cell in state 1, hence concluding the proof. \square

3.4 Glauber dynamics

We will now present a family of stochastic IPS having the property to solve the spacing problem “with an arbitrary precision”, by converging to a distribution on configurations (with same density as the initial configuration) for which the “energy” can be controlled. Let us thus precise this notion of energy.

The energy $E(x)$ of a configuration x is the number of (horizontal or vertical) pairs 11 plus the number of (horizontal or vertical) pairs 00 in the configuration. Precisely, for $x \in \mathcal{E}$, $E(x) = \sum_{(i,j) \in \mathcal{I}_4} \mathbb{1}[x_i = x_j]$, where we recall that \mathcal{I}_4 is the set of pairs of adjacent cells.

Let \mathcal{E}_k be the set of configurations of \mathcal{E} which contain k cells in state 1.

Lemma 1. *If the grid is even-sized, then the configurations of \mathcal{E}_k of minimal energy are exactly the configurations of $\mathcal{E}_k \cap \mathcal{A} = \mathcal{E}_k \cap \mathcal{A}_q$, where q is the minority state.*

of Lemma 1. For $x \in \mathcal{E}$ and $q \in \{0, 1\}$, let us set $E^q(x) = \sum_{(i,j) \in \mathcal{I}_4} \mathbb{1}[x_i = x_j = q]$.

For a configuration $x \in \mathcal{E}_k$, if we count in two different ways the number of pairs $u \in \mathcal{I}_4$ with one cell in state 0 and the other in state 1, we obtain the equality $\text{Card } \mathcal{I}_4 - E(x) = 4k - E^1(x)$. The energy $E(x)$ is thus minimal within

\mathcal{E}_k if and only if $E^1(x)$ is minimal. In particular, if there are configurations $x \in \mathcal{E}_k$ such that $E^1(x) = 0$ (which is equivalent to $x \in \mathcal{A}_1$), then these are exactly the configurations of minimal energy. Similarly, if there are configurations $x \in \mathcal{E}_k$ such that $E^0(x) = 0$ (which is equivalent to $x \in \mathcal{A}_0$), then these are exactly the configurations of minimal energy. Since the grid is even-sized, one can check that whatever the value of k , the set $\mathcal{E}_k \cap \mathcal{A}$ is non-empty. And we have $\mathcal{E}_k \cap \mathcal{A} = \mathcal{E}_k \cap \mathcal{A}_q$, where q is the minority state (or any state in case of equality). \square

For $(i, j) \in \mathcal{I}_4$, let us also define the local energy $E_{(i,j)}(x)$ of configuration x at edge (i, j) by:

$$E_{(i,j)}(x) = \sum_{(k,\ell) \in \mathcal{V}_4(i,j)} \mathbf{1}[x_k = x_\ell],$$

where $\mathcal{V}_4(i, j)$ denotes the six edges of $\mathcal{I}_4 \setminus \{(i, j)\}$ sharing a vertex with (i, j) .

Let $\beta \in \mathbb{R}$ be some fixed parameter. We propose the stochastic IPS dynamics defined as follows.

1. Choose uniformly at random a pair $u = (i, j) \in \mathcal{I}_4$ of horizontal or vertical consecutive cells.
2. Then, exchange the states of cells i and j with probability $p(x, u)$ defined by

$$p(x, u) = \frac{\exp(\beta E_u(x))}{\exp(\beta E_u(x)) + \exp(\beta(6 - E_u(x)))} = \frac{1}{1 + \exp(\beta(6 - 2E_u(x)))}.$$

The number of cells in state 1 is conserved by this dynamics, so that for any $k \in \{0, N_{\mathcal{L}}\}$, it defines a discrete time Markov chain on \mathcal{E}_k . The sequence of edges that will be chosen at each time step is given by a sequence of i.i.d. random variables $(u_t)_{t \in \mathbb{N}} \in \mathcal{I}_4^{\mathbb{N}}$, where u_t is uniformly distributed on \mathcal{I}_4 . Starting from an initial condition $x \in \mathcal{E}_k$, the system evolves according to the sequence of states $(x^t)_{t \in \mathbb{N}}$ defined by $x^0 = x$ and

$$x^{t+1} = \begin{cases} \tau(x^t, u_t) & \text{with probability } p(x, u_t), \\ x^t & \text{with probability } 1 - p(x, u_t). \end{cases}$$

This Markov chain is clearly irreducible and aperiodic. We denote its transition kernel by P . In particular, if $x \neq y$ and $y = \tau(x, u)$ for some $u \in \mathcal{I}_4$, then we have: $P(x, y) = \frac{1}{\text{Card } \mathcal{I}_4} p(x, u)$.

Proposition 5. *The Markov chain defined above is reversible, and its stationary distribution on \mathcal{E}_k is given by $\mu_\beta(x) = \frac{1}{Z_\beta} \exp(-\beta E(x))$ for all $x \in \mathcal{E}_k$, where $Z_\beta = \sum_{x \in \mathcal{E}_k} \exp(-\beta E(x))$.*

Proof. Let us check that the detailed balance $\mu_\beta(x)P(x, y) = \mu_\beta(y)P(y, x)$ holds for any two configurations x and y . It is enough to prove that if $x \neq y$ and $y = \tau(x, u)$ for some $u \in \mathcal{I}_4$, then $\exp(-\beta E(x))p(x, u) = \exp(-\beta E(y))p(y, u)$. But in that case, $E_u(y) = 6 - E_u(x)$ and $E(y) - E_u(y) = E(x) - E_u(x)$, so that the equality is satisfied. \square

Proposition 6. *If the grid is even-sized, then when $\beta \rightarrow \infty$, the distributions μ_β converge to the uniform measure on configurations of minimal energy, that is to the uniform measure on $\mathcal{E}_k \cap \mathcal{A}$.*

Proof. It follows from the definition of μ_β , and Lemma 1 above. \square

To sum up, the Glauber dynamics gives a simple way to approach our goal of spacing out particles. Compared to our checkerboard synchronisation rule Φ_C , it has the advantage of being simple and to use only horizontal and vertical interactions between cells. The distribution at the equilibrium can be determined analytically: it has same weight on all archipelagos with same number of particles, and the weight of non-archipelagos decreases exponentially as a function of β . But, as for Φ_C , there is still the need to know what are the time scales for observing the convergence to the equilibrium: a rule that would converge with a speed that is exponentially slow with the grid size would be useless in practice.

4 The density classification problem on finite lattices

As a first step, we propose to study here only how to use a modified version of Φ_C to solve the density classification problem, leaving the Glauber dynamics for future work. We build our solution to the density classification problem by combining Φ_C with the majority rule.

For the checkerboard synchronisation dynamics, we introduce the following parametric variant of Φ_C , in order to increase the speed of convergence. For a configuration $x \in \mathcal{E}$ and a pair $(i, j) \in \mathcal{I}_4$, the new rule $\tilde{\Phi}_C$ is defined as follows: (a) if both i and j have exactly one adjacent cell in the same state, the exchange (which is then always allowed in Φ_C) is now applied with a probability λ ; (b) in all other cases, we apply Φ_C ; (c) for a pair $(i, j) \in \mathcal{D}_4$, the exchange (which is always allowed in Φ_C), is now done only with a probability χ . For $\lambda = \chi = 1$, we recover Φ_C .

We now combine $\tilde{\Phi}_C$ with a majority rule, to obtain a rule Φ_D defined as follows: for a configuration $x \in \mathcal{E}$ and a pair $(i, j) = u \in \mathcal{I}$,

$$\Phi_D(x, u) = \begin{cases} \tilde{\Phi}_C(x, u) & \text{with probability } 1 - \epsilon, \\ \text{Maj}(x, i) & \text{with probability } \epsilon/2, \\ \text{Maj}(x, j) & \text{with probability } \epsilon/2, \end{cases}$$

where $\text{Maj} : \mathcal{E} \times \mathcal{L} \rightarrow \mathcal{E}$ is the function such that $y = \text{Maj}(x, i)$ is defined by $y_c = x_c$ if $c \neq i$ and y_i is the majority state in the Moore neighbourhood of c (the 8 nearest neighbours of c). We can now state our main proposition.

Proposition 7. *For an even-sized grid $\mathcal{L} = (\mathbb{Z}/a\mathbb{Z}) \times (\mathbb{Z}/b\mathbb{Z})$ with $a, b \in 2\mathbb{Z}$, for any configuration $x \in \mathcal{L}$, and any non-zero value of λ and χ , the probability that Φ_D provides a good classification of x tends to 1 as ϵ tends to 0.*




Figure 4: Time to convergence to an archipelago for different values of L . Left: λ varies with $\chi = 0.1$. Right: χ varies with $\lambda = 0.25$

Sketch. The proof is the same as the one given for a one-dimensional system [5]. For the sake of simplicity we can set $\lambda = 1$ and $\chi = 1$, that is, make $\tilde{\Phi}_C$ and Φ_C equal. First, remark that despite the stochastic nature of our systems, archipelagos are well classified with probability one. In particular, if q is the minority state of x , then the sets of sub-checkerboards C_e^q and C_o^q are stable by the application of both $\tilde{\Phi}_C$ and the majority rule. Moreover, the majority rule either leaves x unchanged or diminishes by 1 the number of q 's in x . In other words, the system can only converge to the right fixed point 0^L or 1^L .

The second property to remark is that as ϵ gets smaller, the probability of *not* applying the majority rule during the first k time steps tends to 1 for every value of k . In other words, for a configuration x with minority symbol q , we can make the probability to reach a sub-checkerboard in $C_e^q \cup C_o^q$ before applying the majority rule as high as needed. \square

In order to evaluate the quality of the rule in practice, let us now briefly explore how the two rules $\tilde{\Phi}_C$ and Φ_D behave with respect to their various settings.

For the particle spacing problem, our rule is defined with three parameters: the grid width L and the two probabilities of exchange λ and χ .

Let us first examine how to set λ and χ . For each setting of the system, we repeated 1000 experiments consisting of initializing the system with an independent Bernoulli of parameter $1/2$ for each cell and measuring the time needed to attain an archipelago configuration. Note that for the sake of making fair comparisons, we show here the *rescaled time*, that is, a time step is taken as L^2 random updates of the global rule. We also take L to be even as for odd-sized grids there are initial conditions for which the particle spacing problem has no solution. In the previous studies, typical values of L were taken around 20, see e.g. Ref. [3, 4] and references therein.




Figure 5: Quality of classification as a function of the grid width L . (Even and odd values are shown on separate curves.) Settings are: $\lambda = 0.25$, $\chi = 0.1$, $\epsilon = 0.001$.

Figure 4-left shows how the average time to convergence varies as a function of λ for the specific value $\chi = 0.1$. It can be observed that for λ smaller than 0.4 the time is relatively small, while for higher values of λ , the time increases drastically.

To examine the effect of χ , we arbitrarily fixed the value of λ to 0.25 and measured the average convergence time to an archipelago. Figure 4-right shows how the average time to convergence varies as a function of χ . Here again, it can be observed that for χ smaller than 0.3 the time is relatively small, while for higher values of χ , the time increases drastically.

Interestingly, these two experiments show that in order to avoid the existence of non-archipelago fixed-point configurations, the two parameters have to be set strictly greater than zero, but can not be set too high. It is an open question to determine if there exists a phase transition with respect to the convergence to an archipelago. This would imply that, for infinite systems, if λ and χ are set above a given threshold, with a high probability, the system does not converge to an archipelago.

Our third experiment is to observe the density classification itself. For each random sample, we took random initial conditions with a uniform probability to be 0 or 1 for each cell independently. For even-sized grids, in case of equality between the number of 0s and 1s, we dismissed this initial condition and re-sampled another one with the same random distribution. We define the *quality* as the ratio of successful classifications, that is, the convergence to right fixed point 0^L or 1^L depending on whether the initial condition has a density smaller than or greater than $1/2$.

Fig. 5 shows the evolution of the quality as a function of the grid width L , for the particular setting $\lambda = 0.25$, $\chi = 0.1$, $\epsilon = 0.001$. (Surprisingly, we empirically remarked that the quality for $\chi = 0.2$ is slightly lower.) These results show that this rule has a quality that is comparable to the best two-dimensional classification rules known so far [3, 4]. Without surprise, for even-sized grids, the

quality decreases as L gets larger, as it becomes more difficult to discriminate between the configurations that have approximately the same number of 0s and 1s. The curve for odd-sized grids is more surprising. Indeed, it shows an *increase* of the quality with L , at least in the range of sizes that were examined. We believe that this phenomenon results from the impossibility to space out the particles for some configurations of odd-sized grid. The system is in some sense “forced” to converge to a non-perfect configuration, in which the majority rule may introduce errors and make the system shift towards the wrong fixed point. However, as the size further increases this effect is less important and it is probable that for a given ϵ , the difference between even-sized and odd-sized grids disappears.

5 Some questions

The density classification problem and the particle spacing problem can both be extended to infinite lattices. The set of configurations is then $\mathcal{E} = \{0, 1\}^{\mathbb{Z}^d}$. For the density classification problem, a possible extension to infinite lattices consists in designing a cellular automaton on \mathcal{E} such that if the initial configuration is drawn independently for each cell according to a Bernoulli law of parameter p , then if $p < 1/2$, the density of 1s converges to 0, while if $p > 1/2$, the density of 0s converges to 0.

This problem has already been studied by Marcovici and her collaborators [2]. In particular, it was shown that there is a simple example of deterministic cellular automaton that classifies the density on \mathbb{Z}^2 : Toom’s rule, which is the majority rule on the neighbourhood $\mathcal{N} = \{(0, 0), (0, 1), (1, 0)\}$. However, in dimension 1, it is an open problem whether there exists a (deterministic or probabilistic) rule that classifies the density. Taati has partially answered this question by giving an argument that holds for densities close to zero or to one [11].

Similarly, for the particle spacing problem, an extension to infinite lattices consists in asking to design a conservative cellular automaton on \mathcal{E} such that if the initial configuration is drawn independently for each cell according to a Bernoulli law of parameter p , then if $p < 1/2$, the density of non-isolated 1s converges to 0, while if $p > 1/2$, the density of non-isolated 0s converges to 0. It is known that the traffic cellular automaton F_{184} is a solution to that problem on \mathbb{Z} , but the problem remains open in dimension $d \geq 2$ [1].

The IPS models we have introduced in Section 3 are also interesting when studying them on \mathbb{Z}^2 instead of finite lattices. In that case, to define properly the model, we need to consider continuous-time updates: each interacting pair of \mathcal{I} possess a clock that rings at times that are exponentially distributed (independently for the different pairs), and the local rule is applied when the clock rings. It is an open problem to know if there is a proper setting of the checkerboard synchronisation dynamics having the property to space particles on \mathbb{Z}^2 . Another interesting model is the Glauber dynamics for $\beta = \infty$. In that case, we allow exchanges between cells only if it makes the energy decrease

(the exchange is made with probability $1/2$ if the exchange does not change the value of the energy). On finite grids, this IPS has many fixed points that are not archipelagos, but starting from a configuration on \mathbb{Z}^2 drawn according to a Bernoulli measure, the behaviour could be different.

As far as the performance of the models is concerned, we can ask what are the best settings to obtain a good trade-off between the quality of classification and the time needed to converge to a fixed point.

We also ask if we can transform our IPS into probabilistic cellular automata for solving the two-dimensional density classification problem. This can be done by using more states or by sharing the randomness of the cells (see Ref. [9] and references therein), but it is an open problem whether there is a solution within the usual framework of binary probabilistic cellular automata.

Acknowledgments

The work of Siamak Taati is supported by ERC Advanced Grant 267356-VARIS of Frank den Hollander.

References

- [1] Vladimir Belitsky and Pablo A. Ferrari. Invariant measures and convergence properties for cellular automaton 184 and related processes. *Journal of Statistical Physics*, 118(3-4):589–623, 2005.
- [2] Ana Bušić, Nazim Fatès, Jean Mairesse, and Irène Marcovici. Density classification on infinite lattices and trees. *Electronic Journal of Probability*, 18(51):1–22, 2013.
- [3] Pedro P. B. de Oliveira. On density determination with cellular automata: Results, constructions and directions. *Journal of Cellular Automata*, 9(5-6):357–385, 2014.
- [4] Nazim Fatès. A note on the density classification problem in two dimensions. In *Proceedings of Automata 2012, the 18th International Workshop on Cellular Automata and Discrete Complex Systems: Exploratory papers*, pages 11–18, 2012. Rapport de recherche I3S/RR-2012-04-FR - Text also available at: <http://hal.inria.fr/hal-00727558>.
- [5] Nazim Fatès. Stochastic cellular automata solutions to the density classification problem - when randomness helps computing. *Theory of Computing Systems*, 53(2):223–242, 2013.
- [6] Henryk Fukś. Solving two-dimensional density classification problem with two probabilistic cellular automata. *Journal of Cellular Automata*, 10(1-2):149–160, 2015.

- [7] Jarkko Kari and Bastien Le Gloannec. Modified traffic cellular automaton for the density classification task. *Fundamenta Informaticae*, 116(1-4):141–156, 2012.
- [8] Mark Land and Richard K. Belew. No perfect two-state cellular automata for density classification exists. *Physical Review Letters*, 74(25):5148–5150, Jun 1995.
- [9] Jean Mairesse and Irène Marcovici. Around probabilistic cellular automata. *Theoretical Computer Science*, 559(0):42–72, 2014.
- [10] Norman H. Packard. *Dynamic Patterns in Complex Systems*, chapter Adaptation toward the edge of chaos, pages 293–301. World Scientific, Singapore, 1988.
- [11] Siamak Taati. Restricted density classification in one dimension. In Jarkko Kari, editor, *Proceedings of AUTOMATA 2015*, volume 9099 of *LNCS*, pages 238–250. Springer, 2015.